

VARIABLE SELECTION AND MODEL AVERAGING IN SEMIPARAMETRIC OVERDISPERSED GENERALIZED LINEAR MODELS

REMY COTTET, ROBERT KOHN, AND DAVID NOTT

Summary

Flexibly modeling the response variance in regression is important for efficient parameter estimation, correct inference, and for understanding the sources of variability in the response. Our article considers flexibly modelling the mean and variance functions within the framework of double exponential regression models, a class of overdispersed generalized linear models. The most general form of our model describes the mean and dispersion parameters in terms of additive functions of the predictors. Each of the additive terms can be either null, linear, or a fully flexible smooth effect. When the dispersion model is null the mean model is linear in the predictors and we obtain a generalized linear model, whereas with a null dispersion model and fully flexible smooth terms in the mean model we obtain a generalized additive model. Whether or not to include predictors, whether or not to model their effects linearly or flexibly, and whether or not to model overdispersion at all is determined from the data using a fully Bayesian approach to inference and model selection. Model selection is accomplished using a hierarchical prior which has many computational and inferential advantages over priors used in previous empirical Bayes approaches to similar problems. We describe an efficient Markov chain Monte Carlo sampling scheme and priors that make the estimation of the model practical with a large number of predictors. The methodology is illustrated using real and simulated data.

Key Words: Bayesian analysis; Double exponential family; Hierarchical priors; Variance estimation.

1. INTRODUCTION

Flexibly modelling the response variance in regression is important for efficient estimation of mean parameters, correct inference and for understanding the sources of variability in the response. Response distributions that are commonly used for modelling non-Gaussian data such as the binomial

and Poisson, although natural and interpretable, have a variance that is a function of the mean and often real data exhibits more variability than might be implied by the mean-variance relationship, a phenomenon referred to as overdispersion. Underdispersion, where the data exhibits less variability than expected, can also occur, although this is less frequent.

Generalized linear models (GLMs) have traditionally been used to model non-Gaussian regression data (Nelder and Wedderburn, 1972, McCullagh and Nelder, 1989), where the response y has a distribution from the exponential family and a transformation of the mean response is a linear function of predictors. This framework is extended to generalized additive models (GAMs) by Hastie and Tibshirani (1990) where a transformation of the mean is modelled as a flexible additive function of the predictors. However, the restriction to the exponential family in GLMs and GAMs is sometimes not general enough. While there is often strong motivation for using exponential family distributions on the grounds of interpretability, the variance of these distributions is a function of the mean and the data often exhibit greater variability than is implied by such mean-variance relationships.

Quasi-likelihood (Wedderburn, 1974) provides one simple approach to inference in the presence of overdispersion, where the exponential family assumption is dropped and only a model for the mean is given with the response variance a function of the mean up to a multiplicative constant. However, this approach does not allow overdispersion to be modelled as a function of covariates. An extension of quasi-likelihood which allows this is the extended quasi-likelihood of Nelder and Pregibon (1987), but in general extended quasi-likelihood estimators may not be consistent (Davidian and Carroll, 1988). Use of a working normal likelihood for estimating mean and variance parameters can also be used for modelling overdispersion (Peck *et al.*, 1984). However, the non-robustness of the estimation of mean parameters when the variance function is incorrectly specified is a difficulty with this approach. Alternatively, a generalized least squares estimating equation for mean parameters can be combined with a normal score estimating equation for variance parameters, a procedure referred to as pseudolikelihood (Davidian and Carroll, 1987). Both the working normal likelihood and pseudolikelihood approaches are related to the theory of generalized estimating equations (see Davidian and Giltinan, 1995, p. 57). Additive extensions of generalized estimating equations are considered by Wild and Yee (1996). Smyth (1989) considers modelling the mean and variance in a parametric class of models which allows normal, inverse Gaussian and gamma response distributions, and a quasi-likelihood extension is also proposed which uses a similar approach to pseudolikelihood for estimation of variance parameters. Smyth and Verbyla (1999) consider extensions of residual maximum likelihood (REML) estimation of variance parameters to double generalized linear models where dispersion parameters are modelled linearly in terms of covariates after transformation by a link function.

Inference about mean and variance functions using estimating equations has the drawback that there is no fully specified model, making it difficult to deal with characteristics of the predictive distribution for a future response, other than its mean and variance. Model based approaches to modelling overdispersion include exponential dispersion models and related approaches (Jorgensen, 1997, Smyth, 1989), the extended Poisson process models of Faddy (1997) and mixture models such as the beta-binomial, negative binomial and generalized linear mixed models (Breslow and Clayton, 1993, Lee and Nelder, 1996). One drawback of mixture models is that they are unable to model underdispersion. Generalized additive mixed models incorporating random effects in GAMs are considered by Lin and Zhang (1999). Both Yee and Wild (1996) and Rigby and Stasinopoulos (2005) consider very general frameworks for additive modelling and algorithms for estimating the additive terms. There is clearly scope for further research on inference: Rigby and Stasinopoulos (2005) suggest that one use for their methods is as an exploratory tool for a subsequent fully Bayesian analysis of the kind considered in our article. See also Brezger and Lang (2005) and Smith and Kohn (1996) for other recent work on Bayesian generalized additive models.

Our framework for flexibly modelling the mean and variance functions is based on the double exponential regression models introduced by Efron (1986), an approach which is also related to the extended quasi-likelihood of Nelder and Pregibon (1987). The double exponential family has been further extended by Gelfand and Dalal (1990) and Gelfand et al. (1997). Double exponential regression models do not suffer the drawbacks of extended quasi-likelihood which occur because the extended quasi-likelihood is not a proper log likelihood function. Semiparametric double exponential regression models can be used to extend both generalized linear models and generalized additive models and are able to model both overdispersion and underdispersion. They provide a convenient framework for modelling as they have the parsimony and interpretability of GLMs, while allowing, if necessary, the flexible dependence of the link transformed mean and variance parameters on predictors. The most general model considered in our article describes the mean and dispersion parameters after transformation by link functions as additive functions of the predictors. For each of the additive terms in the mean and dispersion models we are able to choose between no effect, a linear effect or a fully flexible effect. For a null dispersion model and linear effects in the mean model we obtain generalized linear models, while for a null dispersion model and flexible effects in the mean model we obtain generalized additive models. The main contribution of the paper is to describe a fully Bayesian approach to inference that allows the data to decide whether or not to include predictors, whether or not to model the effect of predictors flexibly or linearly, and whether or not to model overdispersion

at all. We note that an important benefit of variable selection and model averaging is that it produces more efficient model estimates when there are redundant covariates or parameters. As far as we know, alternative approaches to flexible modelling of a mean and variance function do not address similar issues of model selection in a systematic way that is practically feasible when there are many predictors. Nott (2004) considers Bayesian nonparametric estimation of a double exponential family model. However Nott's paper does not consider model averaging, and his priors for the unknown functions and smoothing parameters are very different from those used by our article.

Our article refines and generalizes the work of Shively et al. (1999) and Yau, Kohn and Wood (2003) on nonparametric variable selection and model averaging in probit regression models. These papers use a data-based prior to carry out variable selection. To construct this prior it is necessary to first estimate the model with all flexible terms included, even though the actual fitted model may require only a much smaller number of such terms. This makes the approach impractical when there are a moderate to large number of terms in the full model because the Markov chain simulation breaks down. See Yau et al. (2003), who discuss this problem and also give some strategies to overcome it. Another contribution of our article is to overcome the problems with the data-based prior approach by specifying hierarchical priors for the linear terms and flexible terms in both the mean and variance models. The hierarchical specification in our article is also computationally more efficient than the data-based prior approach because it requires one simulation run through the data, whereas the data-based approach requires two runs, the first to obtain the data-based prior and the second to estimate the parameters of the model. Our approach also applies to variable selection and model averaging in generalized linear models and overdispersed generalized linear models where some or all of the predictors enter the model parametrically. A third contribution of the paper is to develop an efficient Markov chain Monte Carlo (MCMC) sampling scheme for carrying out the computations required for inference in the model.

The paper is organized as follows. Section 2 describes the model, priors and our Bayesian approach to inference and model selection. Section 3 discusses an efficient Markov chain Monte Carlo sampling scheme for carrying out the calculations required for inference. Section 4 applies the methodology to both real and simulated data sets. Section 5 reviews and concludes the paper.

2. MODEL AND PRIOR DISTRIBUTIONS

2.1. The double exponential family. Write the density of a random variable y from a one parameter exponential family as

$$(2.1) \quad p(y; \mu, \phi/A) = \exp \left(\frac{y\psi - b(\psi)}{\phi/A} + c(y, \frac{\phi}{A}) \right),$$

where ψ is a location parameter, ϕ/A is a known scale parameter and $b(\cdot)$ and $c(\cdot, \cdot)$ are known functions. The mean of y is $\mu = b'(\psi)$ and the variance of y is $(\phi/A) b''(\psi)$. This means that $\psi = \psi(\mu)$ is a function of μ and so is the variance. Examples of densities which can be written in this form are the Gaussian, binomial, Poisson and gamma (McCullagh and Nelder, 1989). In (2.1) we write the scale parameter as ϕ/A in anticipation of later discussion of regression models where ϕ is common to all responses but A may vary between responses. A double exponential family is defined from a corresponding one parameter exponential family by

$$(2.2) \quad p(y; \mu, \theta, \phi/A) = Z(\mu, \theta, \phi/A) \theta^{1/2} p(y; \mu, \phi/A)^\theta p(y; y, \phi/A)^{1-\theta},$$

where θ is an additional parameter and $Z(\mu, \theta, \phi/A)$ is a normalizing constant. To get some intuition for this definition consider a Gaussian density with variance 1, and apply the double exponential family construction: the resulting double Gaussian distribution is in fact an ordinary Gaussian density with mean μ and variance $1/\theta$ so that we can think of the parameter θ as a scale parameter modelling overdispersion ($\theta < 1$) or underdispersion ($\theta > 1$) with respect to the original one parameter exponential family density. While the double Gaussian density is simply the ordinary Gaussian density, for distributions like the binomial and Poisson, where the variance is a function of the mean, the corresponding double binomial and double Poisson densities are genuine extensions which allow modelling of the variance. Efron (1986) shows that

$$(2.3) \quad E(y) \approx \mu, \quad \text{Var}(y) \approx \frac{\phi}{A\theta} b''(\psi), \quad \text{and} \quad Z(\mu, \theta, \phi/A) \approx 1$$

with these expressions being exact for $\theta = 1$. Equation (2.3) helps to interpret the parameters in the double exponential model and shows how the GLM mean variance relationship is embedded within the double exponential family, which is important for parsimonious modelling of the variance in regression.

2.2. Semiparametric double exponential regression models. Efron (1986) considers regression models with a response distribution from a double exponential family, such that both the mean parameter μ and the dispersion parameter θ are functions of the predictors. Let y_1, \dots, y_n denote n observed responses, and suppose that μ_i and θ_i are the location and dispersion parameters in the

distribution of y_i . For appropriate link functions $g(\cdot)$ and $h(\cdot)$, we consider the model

$$(2.4) \quad g(\mu_i) = \beta_0^\mu + \sum_{j=1}^p x_{ij} \beta_j^\mu + \sum_{j=1}^p f_j^\mu(x_{ij})$$

$$(2.5) \quad h(\theta_i) = \beta^\theta + \sum_{j=1}^p x_{ij} \beta_j^\theta + \sum_{j=1}^p f_j^\theta(x_{ij}) ,$$

We first discuss equation (2.4) for the mean. This equation has an overall intercept β_0^μ , with the effect of the j th covariate given by the linear term $x_{ij} \beta_j^\mu$ and the nonlinear term $f_j^\mu(x_{ij})$, which is modeled flexibly using a cubic smoothing spline prior. Let $x_{.,j} = (x_{ij}, i = 1, \dots, n)$, for $j = 1, \dots, p$. We standardize each of the covariate vectors $x_{.,j}$ to have mean 0 and variance 1, which makes the $x_{.,j}, j = 1, \dots, p$, orthogonal to the intercept term and comparable in size. This means that the covariates are similar in location and magnitude, which is important for the hierarchical priors used in our article. Making the covariates orthogonal to the intercept diminishes the confounding between the intercept and the covariate. We have also found empirically that the standardization makes the computation numerically more stable.

We now describe the priors on the parameters for the model given by (2.2), (2.4) and (2.5). The prior for β_0^μ is normal but diffuse with zero mean and variance 10^{10} . Let $\beta^\mu = (\beta_1^\mu, \dots, \beta_p^\mu)^T$. To allow the elements of β^μ to be in or out of the model, we define a vector of indicator variables $J^\mu = (J_1^\mu, \dots, J_p^\mu)$ such that $J_l^\mu = 0$ means that β_l^μ is identically 0, and $J_l^\mu = 1$ otherwise. For given J^μ , let β_J^μ be the subvector of nonzero components of β^μ , i.e. those components β_l^μ with $J_l^\mu = 1$. We use the notation $N(a, b)$ for the normal distribution with mean a and variance b , $IG(a, b)$ for the inverse gamma distribution with shape parameter a and scale parameter b and $U(a, b)$ for the uniform distribution on the interval $[a, b]$. With this notation, the prior on β^μ , for a given value of J^μ , is $\beta_J^\mu | J^\mu \sim N(0, b^\mu I)$, where $b^\mu \sim IG(s, t)$, $s = 101$ and $t = 10100$. This choice of parameters produces an inverse Gamma prior with mean 101 and standard deviation 10.15 which worked well across a range of examples, both parametric and nonparametric. However, in general the choice of s and t may depend on the scale and location of the dependent variable and is left to the user. For a continuous response, standardizing the dependent variable may be useful here. The issue of sensitivity to prior hyperparameters is addressed later in the simulations of Section 4.5. We also assume that $Pr(J_l^\mu = 1 | \pi^{\beta^\mu}) = \pi^{\beta^\mu}$ for $l = 1, \dots, p$ and that the J_l are independent given π^{β^μ} . The prior for π^{β^μ} is $U(0, 1)$.

We now specify the priors for the nonlinear terms $f_j^\mu, j = 1, \dots, p$. The discussion below assumes that each $x_{.,j}$ is rescaled to the interval $[0, 1]$ so that the priors in the general case are obtained by transforming back to the original scale. Note that we make the assumption of scaling of the predictors to $[0, 1]$ for expository purposes only to simplify notation in our description of the priors, since we have

previously assumed that our covariates are scaled to have mean zero and variance one. We assume that the functions f_j^μ are a priori independent of each other, and for any m abscissae z_1, \dots, z_m , the vector $(f_j^\mu(z_1), \dots, f_j^\mu(z_m))^T$ is normal with zero mean and with

$$\text{cov}(f_j^\mu(z), f_j^\mu(z')) = \exp(c_j^\mu) \Omega(z, z'),$$

where

$$(2.6) \quad \Omega(z, z') = \frac{1}{2} z^2 (z' - \frac{1}{3} z), \quad \text{for } 0 \leq z \leq z' \leq 1.$$

and $\Omega(z', z) = \Omega(z, z')$. This prior on f leads to a cubic smoothing spline for the posterior mean of f_j^μ (Wahba, 1990, p. 16) with $\exp(c_j^\mu)$ the smoothing parameter.

For $j = 1, \dots, p$, let $f_j^\mu(x_{.,j}) = (f_j^\mu(x_{1,j}), \dots, f_j^\mu(x_{n,j}))^T$, and define the $p \times p$ matrix V_j^μ as having (i, k) th element $\Omega(x_{ij}, x_{kj})$, so that $\text{cov}(f_j^\mu(x_{.,j})) = \exp(c_j^\mu) V_j^\mu$. The matrix V_j^μ is positive definite and can be factored as $V_j^\mu = Q_j^\mu D_j^\mu Q_j^{\mu T}$, where Q_j^μ is an orthogonal matrix of eigenvectors and D_j^μ is a diagonal matrix of eigenvalues. Let $W_j^\mu = Q_j^\mu (D_j^\mu)^{\frac{1}{2}}$. Then $f_j^\mu(x_{.,j}) = W_j^\mu \alpha_j^\mu$, where $\alpha_j^\mu \sim N(0, \exp(c_j^\mu) I)$.

To allow the term f_j^μ to be in or out of the model we introduce the indicator variable K_j^μ so that $K_j^\mu = 0$ means that $\alpha_j^\mu = 0$, which is equivalent to $f_j^\mu = 0$. Otherwise $K_j^\mu = 1$. We also force f_j^μ to be null if the corresponding linear term $\beta_j^\mu = 0$, i.e. if the linear term is zero then we force the flexible term to also be zero. If $J_j^\mu = 1$, then we assume that K_j^μ is 1 with a probability π^{f^μ} , with the prior on π^{f^μ} uniform. When $K_j^\mu = 1$, the prior for c_j^μ is $N(a^{c^\mu}, b^{c^\mu})$, where $a^{c^\mu} \sim N(0, 100)$ and $b^{c^\mu} \sim IG(s, t)$, where s and t are defined above.

As a practical matter, we order the eigenvalues D_j^μ of V_j^μ in decreasing order and set to zero all but the largest m eigenvalues, where m is chosen to be the smallest number such that $\sum_{j=1}^m D_j^\mu / \sum_{j=1}^n D_j^\mu \geq 0.98$. In our work m is usually quite small, around 3 or 4. By setting D_j^μ to zero for $j > m$, we set the corresponding elements of α_j^μ to zero and it is therefore only necessary to work with an α_j^μ that is low dimensional. This achieves a parsimonious parametrization of f_j^μ while retaining its flexibility as a prior. Our approach is similar to the pseudospline approach of Hastie (1996) and has the advantage over other reduced spline basis approaches such as those used in Eilers and Marx (1996), Yau, Kohn and Wood (2003) and Ruppert, Wand and Carroll (2003) of not requiring the choice of the number or location of knots.

The interpretation of equation (2.5) for the variance is similar to that of the mean equation. Let $\beta^\theta = (\beta_0^\theta, \dots, \beta_p^\theta)^T$ and define the indicator variable $J^\theta = 0$ if β^θ is identically zero, with $J^\theta = 1$ otherwise, i.e., in the variance equation (unlike the mean equation) all the linear terms are either in

or out of the model simultaneously so we assume that there is linear over or underdispersion in all the variables or none of them. It would not be difficult to do selection on the linear terms for individual predictors in the variance model, but we feel that in many applications it may be possible that there is no overdispersion, so that the null model where all predictors are excluded from the variance model is inherently interesting, with inclusion of all linear terms with a shrinkage prior on coefficients a reasonable alternative. Our prior parametrizes this comparison directly. When $J^\theta = 1$, we take the prior $\beta^\theta \sim N(0, b^\theta I)$, with $b^\theta \sim IG(s, t)$ where s and t are defined above, and $Pr(J^\theta = 1) = 0.5$.

The hierarchical prior for the nonlinear terms f_j^θ is similar to that for f_j^μ . We write $f_j^\theta(x_{.,j}) = W_j^\theta \alpha_j^\theta$, with $\alpha_j^\theta \sim N(0, \exp(c_j^\theta)I)$. The prior for c_j^θ is $N(a^{c^\theta}, b^{c^\theta})$, with $a^{c^\theta} \sim N(0, 100)$, $b^{c^\theta} \sim IG(s, t)$ where s and t are defined above and K_j^θ is 1 with a probability π^{J^θ} , with the prior on π^{J^θ} uniform. We allow the nonlinear terms to be identically zero by introducing the indicator variables $K_j^\theta, j = 1, \dots, p$, where $K_j^\theta = 1$ means that f_j^θ is in the model and $K_j^\theta = 0$ means that it is not. Similarly to the linear case, we impose that $K_j^\theta = 0$ for all j if $J^\theta = 0$, i.e. if $J^\theta = 0$ then all the nonlinear terms in the variance are 0.

This completes the prior specification. The hierarchical prior is specified in terms of indicator variables that allow selection of linear or flexible effects for variables in the mean and dispersion models. We usually use a log link in the dispersion model, $h(\theta) = \log \theta$, and note that in this case $J^\theta = 0$ implies that all θ values are fixed at one, corresponding to no overdispersion. In some of the examples below we will sometimes fix $J^\theta = 0$ which means that our prior gives a strategy for generalized additive modelling with variable selection and the ability to choose between linear and flexible effects for additive terms.

We note that our framework gives an approach to variable selection and model averaging in GLM's and overdispersed GLM's by fixing $K_j^\mu = K_j^\theta = 0, j = 1, \dots, p$, so that all the terms enter the model parametrically. The first example of Section 4 illustrates the ability of our framework to handle situations where a simple parametric model is appropriate.

3. SAMPLING SCHEME

Let Δ be the set of unknown parameters and latent variables in the model. We use Markov chain Monte Carlo to obtain the posterior distributions of functionals of Δ because in general it is impossible to obtain these distributions analytically. For an introduction to Markov chain Monte Carlo methods see e.g. Liu (2001). The idea of Markov chain Monte Carlo is to construct a Markov chain $\{\Delta^{(m)}; m \geq 0\}$ such that the posterior distribution is the stationary distribution of the chain. By running the chain

a long time from an arbitrary starting value $\Delta^{(0)}$, and after discarding an initial “burn in” sequence of b iterations say, where the distribution of the state is influenced by $\Delta^{(0)}$, we will be able to obtain approximate dependent samples from $p(\Delta|y)$. One estimator of a posterior expectation of interest, $E(h(\Delta)|y)$, is

$$\frac{1}{s} \sum_{i=b+1}^{b+s} h(\Delta^{(i)}),$$

where the first b iterations are discarded in taking the sample path average.

When generating the elements of J^μ it is useful to analytically integrate out $\pi^{\beta\mu}$ in the prior for J^μ to obtain

$$p(J^\mu) = B(1 + \sum_l J_l^\mu, 1 + \sum_l (1 - J_l^\mu))$$

where $B(\cdot, \cdot)$ is the Beta function. Similar remarks apply to K^μ and $\pi^{f\mu}$ and K^θ and $\pi^{f\theta}$. Thus $\pi^{\beta\mu}$, $\pi^{f\mu}$ and $\pi^{f\theta}$ do not appear in the sampling scheme below.

The sampling scheme cycles between different kernels for updating subsets of the parameters to construct the transition kernel for the Markov chain. The kernels for updating the subsets are standard Gibbs and Metropolis-Hastings kernels (see Liu, 2001, for further background) . An update of Δ at a given iteration of our sampling scheme proceeds in the following steps, with further details given in the Appendix:

- (1) Sample β_0^μ .
- (2) For $j = 1, \dots, p$, sample (β_j^μ, J_j^μ) as a block.
- (3) Sample b^μ .
- (4) For $j = 1, \dots, p$, sample $(\alpha_j^\mu, c_j^\mu, K_j^\mu)$ as a block.
- (5) Sample $a^{c\mu}, b^{c\mu}$.
- (6) Sample (β^θ, J^θ) as a block.
- (7) Sample b^θ .
- (8) For $j = 1, \dots, q$, sample $(\alpha_j^\theta, c_j^\theta, K_j^\theta)$ as a block.
- (9) Sample $a^{c\theta}, b^{c\theta}$.

At each step the update of a block of parameters is done conditional on the current values for the remaining parameters.

We briefly discuss one important computational issue arising in implementation of the sampling steps above, namely the calculation of the normalizing constants $Z(\mu, \theta, \phi/A)$ in (2.2). Calculation of this quantity is needed in order to calculate the log-likelihood. Our approach is simply to calculate

$Z(\mu, \theta, \phi/A)$ on a fine grid for μ and θ and to then use an interpolation scheme for values of μ and θ not on the grid. In the case of the double Poisson model we sum the density up to a large truncation point where the contribution of remaining terms to the sum is negligible in order to calculate the normalizing constant at the grid points. This is done offline before beginning the MCMC calculations. For the examples discussed later we used a grid for $g(\mu_i)$ and $h(\theta_i)$ extending from $\log(-50)$ to $\log(50)$ in steps of 2. The summation of the density was truncated at 1000 in calculation of the normalizing constants for the double poisson case. For the double binomial case, we considered the same grid for $g(\mu_i)$ and $h(\theta_i)$. Here we calculate the normalizing constant on this grid for all possible values of the weight ϕ/A – the possible weight values are $1/n_i$, $i = 1, \dots, n$ where n_i is the number of trials for the i th observation. Efron (1986) also describes some asymptotic approximations for the normalizing constant but we do not use these approximations here.

4. EMPIRICAL RESULTS

This section illustrates the application of our methodology in a range of examples, starting in Section 4.1 with a simple parametric example involving overdispersed count data. Section 4.2 considers flexible GAM modelling with variable selection for binary data. Section 4.3 considers the same data as in Section 4.2, but incorporates flexible modelling of interaction terms and variable selection on the interactions. This example illustrates that our methodology can handle very large problems where previously proposed empirical Bayes approaches are infeasible. Section 4.4 considers flexible modelling of overdispersed count data, and Section 4.5 summarizes a simulation study examining the frequentist performance of our method.

4.1. Fully parametric regression. This section illustrates our variable selection methodology in a parametric setting by fitting an overdispersed Poisson model to the pox lesions chick data (<http://www.statsci.org>). We think that it is important to start with consideration of a simple fully parametric example to emphasize that our methodology can be applied with the flexible terms excluded for small data sets where it may only be feasible to fit a simple parametric model. Such small data sets are reasonably common in applications of overdispersed generalized linear models.

The dependent variable is the counts of lesions produced on membranes of chick embryos by viruses of the pox group while the independent variable is the level of dilution of the viral medium. There are 58 observations in this data set. The data is analysed in Breslow (1990) and Podlich, Faddy and Smyth (2004). In our model $g(\mu_i) = \log(\mu_i)$ and $h(\theta_i) = \log(\theta_i)$, with g and h linear functions of the dilution level. Figure 1 plots the fit for the parameter μ and for $\log \theta$ as a function of viral dilution. The

posterior probability of overdispersion (that is, the posterior probability of $J_\theta = 1$) is approximately one strongly suggesting that there is overdispersion, and this is consistent with previous analyses by Breslow (1990) and Podlich, Faddy and Smyth (2004). The results show that overdispersion is increasing with an increase of the viral dilution while the lesions count decreases with dilution.




FIGURE 1. Lesions on chick embryos. Plot of the estimated posterior means of mean and variance parameters as a function of viral dilution together with pointwise 95% credible intervals. Panel (a) plots estimated values of the parameter μ along with the data and panel (b) plots estimated values of $\log \theta$.

Figure 2 shows a plot of the log-likelihood versus iteration number in our MCMC sampling scheme as well as the autocorrelation function of the log-likelihood values based on 2000 iterations with 2000 burn in. These plots show that our sampling scheme converges rapidly and mixes well. Corresponding plots for our other examples (not shown) confirm the excellent properties of our sampling scheme. The 4000 iterations of our sampler took 280 seconds on a machine with 2.8 GHz processor. For all the examples considered in this paper programs implementing our sampler were written in Fortran 90.

4.2. Binary logistic regression. This section considers the Pima Indian diabetes dataset obtained from the UCI repository of machine learning databases (<http://www.ics.uci.edu/MLRepository.html>). The data is analysed by Wahba, Gu, Wang and Chapell (1995). A population of women who were at least 21 years old, of Pima Indian heritage and living near Phoenix, Arizona, was tested for diabetes according to World Health Organization criteria. The data were collected by the US National Institute of Diabetes and Digestive and Kidney Diseases. 724 complete records are used after dropping the




FIGURE 2. Lesions on chick embryos. Plot of the log-likelihood versus iteration and estimated autocorrelation based on 2000 iterations with 2000 burn in.

aberrant cases (as in Yau et al. 2003). The dependent variable is diabetic or not according to WHO criteria, where a positive test is coded as “1”. There are eight covariates: number of pregnancies, plasma glucose concentration in an oral glucose tolerance test, diastolic blood pressure (mm Hg), skin triceps skin fold thickness (mm), 2-Hour serum insulin (μ U/ml), body mass index (weight in kg/(height in m)²), diabetes pedigree function, and age in years.

This section fits a main effects binary logistic regression to the data. In the framework of section 2, we are fixing all θ_i at 1 and so are fitting a generalized linear additive model with $g(\mu_i) = \log(\mu_i/(1 - \mu_i))$ that allows for variable selection and choice between flexible and linear effects for the additive terms. The results are shown in figure 3, with the barplot showing the posterior probabilities of effects for each predictor being null, linear and flexible. The barplot suggests that the number of pregnancies, diastolic blood pressure, skin triceps skin fold thickness and 2-Hour serum insulin do not seem to help predict the occurrence of diabetes when the other covariates are in the model. Figure 3 also shows that plasma glucose concentration has a strong positive linear effect, and body mass index, diabetes pedigree function and age have nonlinear effects.

Our method extends the approach of Yau et al. (2003) to any GAM whereas Yau et al. (2003) rely on the probit link to turn a binary regression into a regression with Gaussian errors. Our approach has several other advantages over Yau et al. (2003) as explained in the introduction and section 4.3. Figure 4 shows the results of applying a variant of the data-based priors approach of Yau *et al.* (2003) to the diabetes data.

Since Yau *et al.* (2003) did not consider logistic regression we need to explain how the data-based priors approach was applied here. First, the full model was fitted (all flexible terms included) with a noninformative but proper prior on the c_j^μ parameters ($IG(s, t)$ with $s = 27$ and $t = 1300$). Linear terms are selected together with flexible terms in the approach of Yau *et al.* (2003). The posterior medians and variances for the c_j^μ (with the variances inflated by a factor of n) in this fit of the full model were then used to set means and variances for a normal prior on the c_j^μ in our variable selection prior similar to Yau *et al.* (2003). The results of Figure 4 are similar to those shown in Figure 3 but as we have already discussed the data based priors approach is not feasible in general for doing selection with a large number of terms. Note that in the barplot we only show posterior probabilities for covariate effects being in or out of the model, since linear and flexible terms are selected together in the approach of Yau *et al.* (2003).

We have also compared our approach to that implemented in the GAMLSS R package of Rigby and Stasinopoulos (2005). We implemented a backward stepwise model selection procedure to select between null, linear and flexible effects for the predictors starting with the model containing all flexible terms. We use their generalized AIC criterion with their penalty parameter $\#$ set to 2 to compare models in the backward stepwise procedure. The final model has flexible terms for Age, Glucose concentration and Body mass index, and a linear term for diabetes function. The fit is shown in Figure 5.

We also report for this example the acceptance rates for our Metropolis-Hastings proposal at step 4 of our sampling scheme in the case where there is currently a flexible term in the model for a predictor and a flexible term is also included in the proposed model. We report acceptance rates for each predictor where there is posterior probability greater than 0.1 of inclusion of a flexible term. The corresponding acceptance rates are 19%, 11% and 38% for the three effects selected. These acceptance rates help to quantify the usefulness of Approximation 2 described in the appendix for constructing the Metropolis-Hastings proposal. Results for this example were based on 4000 iterations of our sampling scheme with 4000 burn in. The 4000 iterations took 1300 seconds on a machine with 2.8 GHz processor.

4.3. High dimensional binary logistic regression. This section extends the model for the Pima Indian dataset to allow for flexible second order interactions. This means that the model potentially has 36 flexible terms, 8 main effects and 28 interactions. The purpose of this section is to show how our class of models can handle interactions and that the hierarchical priors allow variable selection with a large number of terms. This is infeasible with the data-based prior approach of Yau *et al.* (2003), as explained in section 1.

We write the generalization of the mean model (2.4) as

$$g(\mu_i) = \beta_0 + \sum_{j=1}^p x_{ij}\beta_j + \sum_{j=1}^p f_j^M(x_{ij}) + \sum_{j=1}^p \sum_{k < j} f_{jk}^I(x_{ij}, x_{ik}) .$$

We have dropped the superscript μ from the β_j and the f_j because we are dealing with the mean equation only. However, we write the flexible main effects and interactions as f_j^M and f_{jk}^I , where M means main effect and I means an interaction. The prior for the f_j^M is the same as for the flexible main effects in Section 2. For the interaction effects we assume that any collection of $\{f_{jk}^I(x_i, z_i), i = 1, \dots, m\}$ is Gaussian with zero mean and

$$\text{cov}(f_{jk}(x, z), f_{jk}(x', z')) = \exp(c_{jk}^I) \Omega(x, x') \Omega(z, z') ,$$

where $\Omega(z, z')$ is defined by equation (2.6). This gives a covariance kernel for the f_{jk} that is the tensor product of univariate covariance kernels (Gu 2002, section 2.4). Once the covariance matrix for $(f_{jk}^I(x_{ij}, x_{ik}), i = 1, \dots, n)$ is constructed, we factor it to get a parsimonious representation as in Section 2. The smoothing parameters c_{jk}^I have a similar prior to the c_j^μ in Section 2. To allow for variable selection of the flexible main effects, let K_j^M be indicator variables such that $K_j^M = 0$ if f_j^M is null and $K_j^M = 1$ otherwise. The prior for K_j^M is the same as for the K_j^μ in Section 2. To allow variable selection on the flexible interaction terms, let K_{jk}^I be an indicator variable which is 0 if f_{jk}^I is null, and is 1 otherwise. To make the bivariate interactions interpretable, we only allow a flexible interaction between the j th and k th variables if both the flexible main effects are in, i.e., if $K_j^M = 0$ or $K_k^M = 0$ or both, then $K_{jk}^I = 0$. If both K_j^M and K_k^M are 1, then

$$p(K_{jk}^I = 1 | K_j^M, K_k^M, \pi^I) = \pi^I$$

where π^I is uniformly distributed. The generation of the interaction effects parameters $(\alpha_{jk}^I, c_{jk}^I, K_{jk}^I)$ is similar to the generation of the other parameters in the model. First the indicator variable is generated from the prior $p(K_{jk}^I | K_j^M, K_k^M)$. If $K_{jk}^I = 1$ then α_{jk}^I, c_{jk}^I are generated as described in the appendix for the generation of the other parameters, otherwise α_{jk}^I is set to zero.

No interactions were detected when the interaction model were fitted to the data. To test the effectiveness of the methodology at detecting interactions we also generated observations from the estimated main effects model, but added an interaction between Diabetes pedigree function and Age. Writing x and z respectively for these two predictors the interaction term added to our fitted additive model for $\log \mu_i / (1 - \mu_i)$ in the simulation takes the simple multiplicative form xz . Table 4.3 reports the results of the estimation when the interaction model was fitted to the artificial data, and shows that the interaction effect between variables 7 and 8 is detected.

	Covariate							
	1	2	3	4	5	6	7	8
Null	0.88	0.00	0.86	0.81	0.76	0.03	0.00	0.00
Linear	0.08	0.58	0.08	0.05	0.11	0.53	0.00	0.00
1	0.00	0.00	0.00	0.00	0.00	0.01	0.03	0.02
2		0.42	0.01	0.03	0.05	0.13	0.24	0.21
3			0.06	0.02	0.02	0.00	0.02	0.03
4				0.14	0.04	0.03	0.07	0.07
5					0.13	0.03	0.06	0.06
6						0.44	0.023	0.23
7							1.00	1.00
8								1.00

TABLE 1. Simulated diabetes data with interaction. Posterior probabilities of null, linear and flexible main effects and flexible interaction effects. The table is interpreted as follows for covariate 4. The posterior probabilities of a null, linear and flexible main effect are 0.81, 0.05 and 0.14. The posterior probability of a flexible interaction between covariates 3 and 4 is 0.02. Other entries in the table are interpreted similarly.

4.4. Double Binomial model. This example considers a dataset in Moore and Tsiatis (1991) and analyzed by Aerts and Claeskens (1997) using a local beta binomial model. An iron supplement was given to 58 female rats at various dose levels. The rats were then made pregnant and sacrificed after 3 weeks. The litter size and the number of fetuses dead were recorded as well as the hemoglobin levels of the mothers. We fitted a double binomial model to the data to try to explain the proportion of dead foetuses with the level of hemoglobin of the mother and litter size as covariates.

Figure 6 summarizes the estimation results and shows the presence of overdispersion. As usual when dealing with binomial like data the count response is rescaled to be a proportion, so that the parameter μ here is on the scale $[0, 1]$. We have used a logistic link for μ and a log link for θ . The results suggest there is no effect for sample size in the mean model, with some support for either linear or flexible effects for hemoglobin in the mean and variance models and for sample size in the variance model.

Similar to the diabetes example, we report the acceptance rates for our Metropolis-Hastings proposals at steps 4 and 8 of our sampling scheme in the case where there is currently a flexible term in the model for a predictor and a flexible term is also included in the proposed model. We report acceptance rates for each predictor and both the mean and the variance model where there is posterior probability greater than 0.1 of inclusion of a flexible term. The acceptance rates for the mean model are 2.5% for hemoglobin and 2.76% for litter size. For the variance model, no flexible effect was selected. Although the acceptance rates are quite low here, our proposals are still good enough to obtain reasonable mixing. Results for this example were based on 5000 iterations of our sampling scheme with 5000 burn in. The 5000 iterations took 4039 seconds on a machine with 2.8 GHz processor.

For this example we also compare an implementation of our methodology using a beta-binomial response distribution to flexible beta-binomial regression implemented in the GAMLSS library in R (Rigby and Stasinopoulos, 2005). Implementation of our method for the beta-binomial family rather than the double exponential is straightforward as our computational scheme makes no particular use of the double exponential family assumption, but only the idea of mean and variance parameters being modelled flexibly as a function of covariates. For beta-binomial regression, Rigby and Stasinopoulos (2005) parametrize the model in terms of a mean parameter μ and dispersion parameter σ which is $\rho/(1-\rho)$ where ρ is the intraclass correlation (if we regard each count observation as an observation of a sequence of exchangeable binary random variables, the intraclass correlation is just the correlation between a pair of these binary random variables). Large value of σ correspond to overdispersion, whereas $\sigma = 0$ corresponds to no overdispersion. Our model is similar to before, except that we replace our model for $h(\theta_i)$ in (2.5) with a model of the same form for $h(\sigma_i)$ where σ_i is the dispersion parameter for observation i and $h(\cdot)$ is a link function which we choose here as the log function.

Figures 7 and 8 show the results of our fit and the GAMLSS fit (with all terms flexible) for the rat data. We can see that the fits are similar.

One advantage of our approach is greater computational stability, a feature that we believe is related to our shrinkage priors. We simulated several datasets from our fitted model for the mean, but assuming no overdispersion ($\sigma = 0$) and then attempted to fit to this simulated data using GAMLSS

and our Bayesian approach with a beta-binomial model. The Bayesian approach produces satisfactory results, but attempting to fit the model in GAMLSS even with only an intercept and no covariates in the variance model results in convergence problems that are not easily resolved (D.M. Stasinopoulos, personal communication). However, the GAMLSS fit is faster, and we have found the GAMLSS package to be very useful for the exploratory examination of many potential models.

4.5. Simulation studies. We consider three simulation studies which show the effectiveness of our methodology for detecting overdispersion when it exists and for distinguishing between null, linear and flexible effects. We also examine the gain in performance which results when our hierarchical variable selection priors are used instead of a similar hierarchical prior in which variable selection is not done. Performance here is measured by the percentage increase in average Kullback-Leibler divergence between the true and estimated predictive densities averaging over observed predictors.

For our fitted overdispersed model to the rats data, we simulated 50 replicate datasets from the fitted model. Table 2 shows the 25th and 75th percentiles of the probabilities of null, linear and flexible effects for the two predictors in the mean and dispersion models over the 50 replications. The results are consistent with our fit to the original data, with appreciable probabilities of linear and flexible effects for sample size in the mean model and sample size and hemoglobin level in the variance model and an appreciable probability for no effect for sample size in the mean model. Table 3 is similar to Table 2 but for the data based priors method of Yau *et al.* (2003), implemented using a similar approach to that discussed in Section 4.2. The results for the data based priors approach are similar to those for our hierarchical priors for the mean model. Posterior probabilities of null effects in the variance model differ in the two implementations, but this may be due to the fact that selection is done separately on different predictors in the variance model in the data based priors approach, whereas we include or exclude predictors together in our hierarchical prior with a shrinkage prior on coefficients. Note that the results of Table 3 show only posterior probabilities for flexible and null effects, as flexible and linear terms are selected together in the approach of Yau *et al.* (2003). Tables 4 and 5 show the results of the simulation study for our method where the hyperparameters s and t in the inverse gamma priors of Section 2.2 are $(s, t) = (6, 500)$ and $(s, t) = (27, 1300)$ (giving prior means of 100 and 50 respectively and standard deviations of 50 and 10 respectively). As can be seen from the tables, the results of our approach are not particularly sensitive to the choice of s and t .

Table 6 is similar to Table 2, but for 50 replicate datasets simulated from a fitted binomial model (that is, with no overdispersion). The probability of a null effect for both covariates in the variance model is near one and again there is high probability of a null effect for sample size in the mean model

Effect	Percentile	Covariate			
		Mean		Var	
		Hem	SS	Hem	SS
Flexible	25 th	0.42	0.15	0.38	0.38
	75 th	0.47	0.27	0.45	0.45
Linear	25 th	0.52	0.19	0.55	0.55
	75 th	0.58	0.30	0.61	0.61
Null	25 th	0.00	0.44	0.00	0.00
	75 th	0.00	0.64	0.00	0.00

TABLE 2. Rats data simulated from the fitted model. The 25th and 75th percentiles of the probabilities of flexible, linear and null effects for the mean and variance components, for the two covariates hemoglobin (Hem) and sample size (SS). Based on 50 replications

Effect	Percentile	Covariate			
		Mean		Var	
		Hem	SS	Hem	SS
Flexible	25 th	1.00	0.40	0.25	0.18
	75 th	1.00	0.48	0.36	0.25
Null	25 th	0.00	0.52	0.64	0.75
	75 th	0.00	0.60	0.75	0.82

TABLE 3. Rats data simulated from the fitted model. The 25th and 75th percentiles of the posterior probabilities of flexible and null effects for the mean and variance components, for the two covariates hemoglobin (Hem) and sample size (SS) for the data based priors approach. Based on 50 replications.

and appreciable probabilities for linear and flexible effects for hemoglobin in the mean model. Table 7 is similar to Table 6 except for the data based priors approach of Yau *et al.* (2003). The results are again similar.

Similarly to Table 6 we examine in Tables 8 and 9 the performance of our approach but with the hyperparameter settings $(s, t) = (6, 500)$ and $(s, t) = (27, 1300)$. Again we see that the results are not sensitive to the hyperparameter settings.

Effect	Percentile	Covariate			
		Mean		Var	
		Hem	SS	Hem	SS
Flexible	25 th	0.10	0.04	0.11	0.10
	75 th	0.13	0.06	0.13	0.13
Linear	25 th	0.86	0.27	0.87	0.88
	75 th	0.89	0.33	0.89	0.90
Null	25 th	0.00	0.60	0.00	0.00
	75 th	0.01	0.69	0.00	0.00

TABLE 4. Rats data simulated from the fitted model. The 25th and 75th percentiles of the probabilities of flexible, linear and null effects for the mean and variance components, for the two covariates hemoglobin (Hem) and sample size (SS). Based on 50 replications and hyperparameter settings $(s, t) = (27, 1300)$.

Effect	Percentile	Covariate			
		Mean		Var	
		Hem	SS	Hem	SS
Flexible	25 th	0.09	0.03	0.11	0.10
	75 th	0.12	0.05	0.14	0.13
Linear	25 th	0.86	0.23	0.86	0.87
	75 th	0.90	0.28	0.89	0.90
Null	25 th	0.00	0.66	0.00	0.00
	75 th	0.02	0.73	0.00	0.00

TABLE 5. Rats data simulated from the fitted model. The 25th and 75th percentiles of the probabilities of flexible, linear and null effects for the mean and variance components, for the two covariates hemoglobin (Hem) and sample size (SS). Based on 50 replications and hyperparameter settings $(s, t) = (6, 500)$.

Table 10 shows probabilities of null, linear and flexible effects for the eight covariates in the diabetes example for 50 simulated replicate datasets from an additive model fitted to the real data. Again the results are consistent with out fit to the full model, with high probability of a null effect for covariates 1, 3, 4 and 5 (the number of pregnancies, diastolic blood pressure, skin triceps skin fold thickness and

Effect	Percentile	Covariate			
		Mean		Var	
		Hem	SS	Hem	SS
Flexible	25 th	0.40	0.08	0.00	0.00
	75 th	0.43	0.22	0.00	0.00
Linear	25 th	0.57	0.12	0.00	0.00
	75 th	0.60	0.36	0.00	0.00
Null	25 th	0.00	0.43	0.99	0.99
	75 th	0.00	0.79	0.99	0.99

TABLE 6. Rats simulated data from a fitted binomial model with no overdispersion. The 25th and 75th percentiles of the probabilities of flexible, linear and null effects are given for the mean and variance components, for the two covariates hemoglobin (Hem) and sample size (SS). Based on 50 replications.

Effect	Percentile	Covariate			
		Mean		Var	
		Hem	SS	Hem	SS
Flexible	25 th	1.00	0.47	0.00	0.00
	75 th	1.00	0.51	0.00	0.00
Null	25 th	0.00	0.49	1.00	1.00
	75 th	0.00	0.53	1.00	1.00

TABLE 7. Rats simulated data from a fitted binomial model with no overdispersion. The 25th and 75th percentiles of the posterior probabilities of flexible and null effects are given for the mean and variance components, for the two covariates hemoglobin (Hem) and sample size (SS), and the data based priors approach. Based on 50 replications.

2-Hour serum insulin respectively), an appreciable probability of a linear effect for covariate 2 (plasma glucose concentration) and high probabilities of nonlinear effects for covariates 6, 7 and 8 (body mass index, diabetes pedigree function and age). Table 11 is similar to 10 but for the data based priors approach of Yau *et al.* (2003). Again the results are similar to those obtained using our hierarchical priors.

Effect	Percentile	Covariate			
		Mean		Var	
		Hem	SS	Hem	SS
Flexible	25 th	0.10	0.02	0.00	0.00
	75 th	0.14	0.05	0.02	0.01
Linear	25 th	0.86	0.10	0.00	0.00
	75 th	0.90	0.19	0.12	0.13
Null	25 th	0.00	0.75	0.86	0.86
	75 th	0.00	0.87	1.00	1.00

TABLE 8. Rats simulated data from a fitted binomial model with no overdispersion. The 25th and 75th percentiles of the probabilities of flexible, linear and null effects are given for the mean and variance components, for the two covariates hemoglobin (Hem) and sample size (SS). Based on 50 replications and hyperparameter settings $(s, t) = (6, 500)$.

Effect	Percentile	Covariate			
		Mean		Var	
		Hem	SS	Hem	SS
Flexible	25 th	0.10	0.02	0.00	0.00
	75 th	0.14	0.05	0.03	0.03
Linear	25 th	0.86	0.14	0.01	0.01
	75 th	0.90	0.25	0.23	0.24
Null	25 th	0.00	0.67	0.73	0.73
	75 th	0.00	0.83	0.99	0.99

TABLE 9. Rats simulated data from a fitted binomial model with no overdispersion. The 25th and 75th percentiles of the probabilities of flexible, linear and null effects are given for the mean and variance components, for the two covariates hemoglobin (Hem) and sample size (SS). Based on 50 replications and hyperparameter settings $(s, t) = (27, 1300)$.

Similar to Table 10 we examine in Tables 12 and 13 the performance of our approach but with the hyperparameter settings $(s, t) = (6, 500)$ and $(s, t) = (27, 1300)$. Again we see that the results are not particularly sensitive to the hyperparameter settings.

Effect	Perc.	Covariate							
		1	2	3	4	5	6	7	8
Flexible	25 th	0.03	0.55	0.06	0.05	0.04	0.62	0.35	0.99
	75 th	0.06	0.61	0.11	0.11	0.12	0.97	0.54	1.00
Linear	25 th	0.03	0.39	0.04	0.03	0.04	0.03	0.19	0.00
	75 th	0.05	0.45	0.08	0.09	0.08	0.38	0.40	0.01
Null	25 th	0.88	0.00	0.80	0.79	0.82	0.00	0.04	0.00
	75 th	0.94	0.00	0.90	0.91	0.92	0.00	0.46	0.00

TABLE 10. Simulated data from the fitted diabetes model. The 25th and 75th percentiles of the posterior probabilities of flexible, linear and null effects are given. Based on 50 replications.

Effect	Perc.	Covariate							
		1	2	3	4	5	6	7	8
Flexible	25 th	0.36	1.00	0.35	0.34	0.28	1.00	0.26	1.00
	75 th	0.52	1.00	0.49	0.49	0.48	1.00	1.00	1.00
Null	25 th	0.48	0.00	0.51	0.51	0.52	0.00	0.00	0.00
	75 th	0.64	0.00	0.65	0.66	0.72	0.00	0.74	0.00

TABLE 11. Simulated data from the fitted diabetes model. The 25th and 75th percentiles of the posterior probabilities of flexible and null effects are given for the data based priors approach. Based on 50 replications.

We now compare the performance of our hierarchical variable selection priors with the same prior but where all terms are flexible (that is, no variable selection is carried out). Our measure of performance is the Kullback-Leibler divergence, averaged over the observed covariates.

In estimating the true response distribution $p_0(y|x)$ using an estimate $\hat{p}(y|x)$ where x denotes the covariates, the Kullback-Leibler divergence is defined as

$$KL(\hat{p}(\cdot|x), p_0(\cdot|x)) = \int p_0(y|x) \log \left[\frac{\hat{p}(y|x)}{p_0(y|x)} \right] dy.$$

We define the average Kullback-Leibler divergence as

$$AKLD(\hat{p}, p_0) = \frac{1}{n} \sum_{i=1}^n KL(\hat{p}(\cdot|x_i), p_0(\cdot|x_i))$$

Effect	Perc.	Covariate							
		1	2	3	4	5	6	7	8
Flexible	25 th	0.04	0.51	0.03	0.06	0.03	0.99	0.42	1.00
	75 th	0.21	0.64	0.20	0.14	0.14	1.00	0.62	1.00
Linear	25 th	0.04	0.36	0.03	0.05	0.04	0.00	0.25	0.00
	75 th	0.10	0.49	0.11	0.10	0.10	0.01	0.41	0.00
Null	25 th	0.67	0.00	0.69	0.77	0.78	0.00	0.00	0.00
	75 th	0.92	0.00	0.94	0.89	0.93	0.00	0.19	0.00

TABLE 12. Simulated data from the fitted diabetes model. The 25th and 75th percentiles of the posterior probabilities of flexible, linear and null effects are given. Based on 50 replications and the hyperparameter settings $(s, t) = (6, 500)$.

Effect	Perc.	Covariate							
		1	2	3	4	5	6	7	8
Flexible	25 th	0.06	0.54	0.07	0.07	0.07	0.99	0.48	0.98
	75 th	0.18	0.69	0.14	0.17	0.14	1.00	0.67	1.00
Linear	25 th	0.04	0.31	0.06	0.04	0.06	0.00	0.17	0.00
	75 th	0.09	0.46	0.11	0.10	0.09	0.01	0.41	0.00
Null	25 th	0.72	0.00	0.75	0.72	0.78	0.00	0.00	0.00
	75 th	0.89	0.00	0.86	0.88	0.86	0.00	0.17	0.01

TABLE 13. Simulated data from the fitted diabetes model. The 25th and 75th percentiles of the posterior probabilities of flexible, linear and null effects are given. Based on 50 replications and the hyperparameter settings $(s, t) = (27, 1300)$.

where x_i , $i = 1, \dots, n$ denotes the observed predictors. Writing $\hat{p}^V(y|x)$ for the estimated predictive density at x for the variable selection prior and $\hat{p}^{NV}(y|x)$ for the estimated predictive density at x for the prior without variable selection we define the average percentage increase in Kullback-Leibler loss for variable selection compared to no variable selection as

$$APKL = \frac{AKLD(\hat{p}^{NV}, p_0) - AKLD(\hat{p}^V, p_0)}{AKLD(\hat{p}^V, p_0)}.$$

When $APKL$ is positive, the prior that allows for variable selection outperforms the prior that does not allow variable selection. Table 14 shows the 10th, 25th, 50th, 75th and 90th percentiles of the

$APKL$ for the 50 replicate data sets generated in our simulation study for the diabetes data, the rats data when no overdispersion is present and the rats data when overdispersion is present. The table shows that the median percentage increase in $APKL$ is positive for all three cases, indicating an improvement for using our hierarchical variable selection prior compared to not doing variable selection. Furthermore, for the rats data with no overdispersion, even the 10th percentile exceeds 28%.

Dataset	Percentiles				
	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
Diabetes	-14.49	-1.17	16.74	40.85	64.14
Rats with overdispersion	-13.92	-5.18	9.08	30.48	62.20
Rats with no overdispersion	28.67	67.55	155.52	293.43	734.54

TABLE 14. 10^{th} , 25^{th} , 50^{th} , 75^{th} and 90^{th} percentiles in the percentage increase in Kullback-leibler divergence when no variable selection is carried out compared to when variable selection is carried out. The results are based on 50 replications for the diabetes data, rats data when no overdispersion is present and rats data when overdispersion is present.

5. CONCLUSION

The article develops a general Bayesian framework for variable selection and model averaging in generalized linear models that allows for over or under dispersion. The priors and sampling are innovative and the flexibility of the approach is demonstrated using a number of examples, ranging from fully parametric to fully nonparametric.

There are a number of natural extensions to the work described here. Although we have implemented our approach to flexible regression for the mean and variance using the double exponential family of distributions, it is easy to implement a similar approach using other distributional families for overdispersed count data such as the beta-binomial and negative binomial. We have demonstrated use of the beta-binomial in one of our real data examples. Flexible modelling of multivariate data could also be easily accommodated in our framework by incorporation of other kinds of random effects apart from those involved in our nonparametric functional forms. These and other extensions are the subject of ongoing research.

ACKNOWLEDGEMENTS

This work was supported by an Australian Research Council Grant. We thank Dr Mikis Stasinopoulos for a quick and helpful response to some questions about the GAMLSS package.

6. APPENDIX

This section gives details of the sampling scheme in Section 3. Most of the steps involve an application of the Metropolis-Hastings method based on one of the following two approximations to the conditional densities.

Approximation 1

Here we seek to generate a parameter ψ from its full conditional $p(\psi|y, \Delta \setminus \psi)$, where Δ consists of all the parameters and latent variables used in the sampling scheme and $\Delta \setminus \psi$ means all of Δ excluding ψ . We write

$$(6.1) \quad \begin{aligned} p(\psi|y, \Delta \setminus \psi) &\propto p(y|\psi, \Delta \setminus \psi)p(\psi|\Delta \setminus \psi) \\ &\propto \exp(-l(\psi)) \end{aligned}$$

where $l(\psi)$ is the negative of the logarithm of the left side of (6.1) and for convenience the dependence of $l(\cdot)$ on $\Delta \setminus \psi$ is omitted.

Let $\hat{\psi}$ be the minimum of $l(\psi)$ and $\Psi = \partial^2 l(\hat{\psi}) / \partial \psi \partial \psi$. We approximate the full conditional of ψ by a normal density with mean $\hat{\psi}$ and covariance matrix Ψ^{-1} . Generally we find $\hat{\psi}$ using numerical optimization routines from the NAG or IMSL libraries and have not experienced any difficulties of convergence. As a practical matter, there may be a substantial benefit to early stopping of the optimization in constructing our proposals after just a few or even one step. Just a few steps gives a good approximation to $\hat{\psi}$ and this suffices to obtain good proposals with the early stopping resulting in a considerable saving in computation time since several applications of Approximation 1 are done at every iteration of the sampling scheme.

Approximation 2

Here we seek to generate parameters ψ and w as a block from their joint conditional density. We assume that $p(y|\psi, w, \Delta \setminus \{\psi, w\}) = p(y|\psi, \Delta \setminus \{\psi, w\})$ and $p(\psi|w, \Delta \setminus \{\psi, w\}) = p(\psi|w)$ is Gaussian.

Then,

$$\begin{aligned} p(w|y, \Delta \setminus \{\psi, w\}) &= \int p(\psi, w|y, \Delta \setminus \{\psi, w\}) d\psi \\ &\propto \int p(y|\psi, \Delta \setminus \{\psi, w\}) p(\psi|w) d\psi \times p(w|\Delta \setminus \{\psi, w\}). \end{aligned}$$

Let $q_1(\psi)$ be a Gaussian approximation to $p(y|\psi, \Delta \setminus \{\psi, w\})$ as in approximation 1, and note that it is independent of w . Let

$$q_2(w) = \int q_1(\psi) p(\psi|w) d\psi$$

and note that $q_2(w)$ can be evaluated explicitly as a function of w . Precisely, if $q_1(\psi)$ is a Gaussian $N(\mu_1, \Sigma_1)$ and $p(\psi|w)$ is a Gaussian $N(\mu_w, \Sigma_w)$ then we obtain

$$\begin{aligned} q_2(w) \propto |\Sigma_w^{-1}|^{-1/2} |\Sigma_1^{-1} + \Sigma_w^{-1}|^{-1/2} \exp \left(-\frac{1}{2} \{ \mu_w^T \Sigma_w^{-1} \mu_w + \right. \\ \left. (\mu_1^T \Sigma_1^{-1} + \mu_w^T \Sigma_w^{-1})^T (\Sigma_1^{-1} + \Sigma_w^{-1})^{-1} (\mu_1^T \Sigma_1^{-1} + \mu_w^T \Sigma_w^{-1}) \} \right). \end{aligned}$$

Thus, we approximate $p(w|y, \Delta \setminus \{\psi, w\})$ by $q_3(w) = q_2(w) p(w|\Delta \setminus \{\psi, w\})$. In our applications w is scalar and so it is straightforward to approximate $q_3(w)$ by a Gaussian as in the first approximation and hence generate w . Once w is generated, it is straightforward to generate ψ from $p(\psi|w)$ which is Gaussian.

In step 1, let $\psi = \beta_0^\mu$ and construct a proposal density for ψ as in the first approximation. We either accept the proposed value or retain the current value according to the usual Metropolis-Hastings rule.

In Step 2, J_j^μ and β_j^μ are generated as a block. If $K_j^\mu = 1$ then $J_j^\mu = 1$. If $K_j^\mu = 0$ then J_j^μ is generated as 0 or 1 from the prior. If J_j^μ is generated as a 1, then β_j^μ is generated from a normal approximation to its full conditional density as in the first approximation. The proposed pair is either accepted or it is rejected in favor of the current values according to the Metropolis-Hastings rule.

In step 3, b^μ is sampled from its full conditional density,

$$(6.2) \quad p(b^\mu|y, \Delta \setminus b^\mu) \propto p(\beta^\mu|b^\mu) p(b^\mu)$$

the right side of (6.2) is the unnormalized inverse gamma density with shape parameter $s + \sum_j J_j^\mu/2$ and scale parameter $t + \beta_j^{\mu T} \beta_j^\mu/2$.

To describe step 4 of the sampling scheme, we first show how c_j^μ and α_j^μ are generated if $K_j^\mu = 1$. Let $\psi = \alpha_j^\mu$ and $w = c_j^\mu$. Then $p(y|\psi, w, \Delta \setminus \{\psi, w\}) = p(y|\psi, \Delta \setminus \{\psi, w\})$ and $p(\psi|w)$ is Gaussian in ψ . We generate (ψ, w) as a block as in the second approximation.

We generate a proposal for $(K_j^\mu, c_j^\mu, \alpha_j^\mu)$ as follows. First, generate a proposal for K_j^μ as 0 if $J_j^\mu = 0$, and 0 or 1 from the prior if $J_j^\mu = 1$. Next, if the proposed value of $K_j^\mu = 1$ then generate (c_j^μ, α_j^μ) as

outlined above. Then, we accept or reject the block proposal according to the Metropolis-Hastings rule.

In step 5, $a^{c\mu}$ and $b^{c\mu}$ can be updated from their full conditional distributions. We have

$$p(a^{c\mu}|\Delta \setminus a^{c\mu}) \propto p(a^{c\mu})p(c^\mu|a^{c\mu}, b^{c\mu})$$

and we recognize the right hand side as an unnormalized normal density. We have

$$p(a^{c\mu}|\Delta \setminus \{a^{c\mu}\}) = N \left(\left(\frac{\sum_j K_j^\mu}{b^{c\mu}} + \frac{1}{100} \right)^{-1} \frac{(c^\mu)^T c^\mu}{b^{c\mu}}, \left(\frac{\sum_j K_j^\mu}{b^{c\mu}} + \frac{1}{100} \right)^{-1} \right).$$

Also,

$$p(b^{c\mu}|\Delta \setminus \{b^{c\mu}\}) \propto p(b^{c\mu})p(c^\mu|a^{c\mu}, b^{c\mu})$$

and we recognize the right hand side as an unnormalized inverse gamma density. We have

$$p(b^{c\mu}|\Delta \setminus \{b^{c\mu}\}) = IG \left(s + \frac{\sum_j K_j^\mu}{2}, t + \frac{1}{2}(c^\mu - a^{c\mu})^T (c^\mu - a^{c\mu}) \right).$$

Step 6 is similar to step 1, except that a multivariate normal approximation is used to generate β^θ given $J^\theta = 1$ and the current values of other parameters. Step 7 is similar to step 3 and

$$p(b^\theta|\Delta \setminus \{b^\theta, J^\theta\}, J^\theta = 1) = IG \left(s + \frac{q+1}{2}, t + \frac{1}{2}(\beta^\theta)^T \beta^\theta \right).$$

Steps 8 and 9 are performed similarly to steps 4 and 5.

REFERENCES

- Aerts, M. and Claskens, G. (1997), “Local polynomial estimation in multiparameter models,” *Journal of the American Statistical Association*, 92, 1536–1545.
- Breslow, N. (1990), “Further studies in the variability of pock counts,” *Statistics in Medicine*, 9, 615–626.
- Breslow, N. and Clayton, D. (1993), “Approximate inference in generalized linear mixed models,” *Journal of the American Statistical Association*, 88, 9–25.
- Brezger, A. and Lang, S. (2005), “Generalized additive structured regression based on Bayesian P-splines,” *Computational Statistics and Data Analysis*, 50, 967–991.
- Davidian, M. and Carroll, R. (1987), “Variance function estimation,” *Journal of the American Statistical Association*, 82, 1079–1091.
- (1988), “A note on extended quasilikelihood,” *Journal of the Royal Statistical Society B*, 50, 74–82.
- Davidian, M. and Giltinan, D. (1995), *Nonlinear Models for Repeated Measurement Data*, New York: Chapman and Hall.

- Efron, B. (1986), "Double exponential families and their use in generalised linear regression," *Journal of the American Statistical Association*, 81, 709–721.
- Eilers, P. H. C. and Marx, B. D. (1996), "Flexible Smoothing with B-splines and Penalties with Rejoinder," *Statistical Science*, 11, 89–121.
- Faddy, M. (1997), "Extended Poisson process modelling and analysis of count data," *Biometrical Journal*, 39, 431–440.
- Gelfand, A. and Dalal, S. (1990), "A note on overdispersed exponential families," *Biometrika*, 77, 55–64.
- Gelfand, A., Dey, D., and Peng, F. (1997), "Overdispersed generalized linear models," *Journal of Statistical Planning and Inference*, 64, 93–107.
- Gu, C. (2002), *Smoothing spline ANOVA models*, New York: Springer-Verlag.
- Hastie, T. (1996), "Pseudosplines," *Journal of the Royal Statistical Society B*, 58, 379–396.
- Hastie, T. and Tibshirani, R. (1990), *Generalized Additive Models*, New York: Chapman and Hall.
- Jorgensen, B. (1997), *The Theory of Dispersion Models*, London: Chapman and Hall.
- Lee, Y. and Nelder, J. (1996), "Hierarchical generalized linear models (with Discussion)," *Journal of the Royal Statistical Society B*, 58, 619–678.
- Lin, X. and Zhang, D. (1999), "Inference in generalized additive mixed models by using smoothing splines," *Journal of the Royal Statistical Society B*, 61, 381–400.
- Liu, J. (2001), *Monte Carlo Strategies in Scientific Computing*, New York: Springer-Verlag.
- McCullagh, P. and Nelder, J. (1989), *Generalized Linear Models*, London: Chapman and Hall, 2nd ed.
- Moore, D. and Tsiatis, A. (1991), "Robust estimation of the variance in moment methods for extra-binomial and extra-Poisson variation," *Biometrics*, 47, 383–401.
- Nelder, J. and Pregibon, D. (1987), "An extended quasi-likelihood function," *Biometrika*, 74, 221–232.
- Nelder, J. and Wedderburn, R. (1972), "Generalized linear models," *Journal of the Royal Statistical Society B*, 34, 370–384.
- Nott, D. (2004), "Semiparametric estimation of mean and variance functions for non-Gaussian data," Working paper.
- Peck, C., Beal, S., Sheiner, L., and Nichols, A. (1984), "Extended least squares nonlinear regression: A possible solution to the choice of weights problem in analysis of individual pharmacokinetic data," *Journal of Pharmacokinetics and Biopharmaceutics*, 12, 545–558.
- Podlich, H., Faddy, M., and Smyth, G. (2004), "Semi-parametric extended Poisson process models," *Statistics and Computing*, 14, 311–321.

- Rigby, R. and Stasinopoulos, D. (2005), “Generalized additive models for location, scale and shape,” *Applied Statistics*, 54, 1–38.
- Ruppert, D., Wand, M., and Carroll, R. (2003), *Semiparametric regression*, Cambridge: Cambridge University Press.
- Shively, S., Kohn, R., and Wood, S. (1999), “Variable selection and function estimation in additive nonparametric regression using a data based prior (with discussion),” *Journal of the American Statistical Association*, 94, 777–807.
- Smith, M. and Kohn, R. (1996), “Nonparametric Regression Using Bayesian Variable Selection,” *Journal of Econometrics*, 75, 317–344.
- Smyth, G. (1989), “Generalized linear models with varying dispersion,” *Journal of the Royal Statistical Society B*, 51, 47–60.
- Smyth, G. K. and Verbyla, A. P. (1999), “Adjusted likelihood methods for modelling dispersion in generalized linear models,” *Environmetrics*, 10, 696–709.
- Wahba, G., Wang, Y., Gu, C., Klein, R., and Klein, B. (1995), “Smoothing spline ANOVA for exponential families, with application to the Wisconsin Epidemiological Study of Diabetic Retinopathy,” *Annals of Statistics*, 23, 1865–1895.
- Wedderburn, R. (1974), “Quasi-likelihoods functions, generalized linear models and the Gauss-Newton method,” *Biometrika*, 61, 439–447.
- Wild, C. and Yee, T. (1996), “Additive extensions to generalized estimating equation methods,” *Journal of the Royal Statistical Society B*, 58, 711–725.
- Yau, P., Kohn, R., and Wood, S. (2003), “Bayesian variable selection and model averaging in high-dimensional multinomial nonparametric regression,” *Journal of Computational and Graphical Statistics*, 12, 23–54.
- Yee, T. and Wild, C. (1996), “Vector Generalized Additive Models,” *Journal of the Royal Statistical Society B*, 58, 481–493.

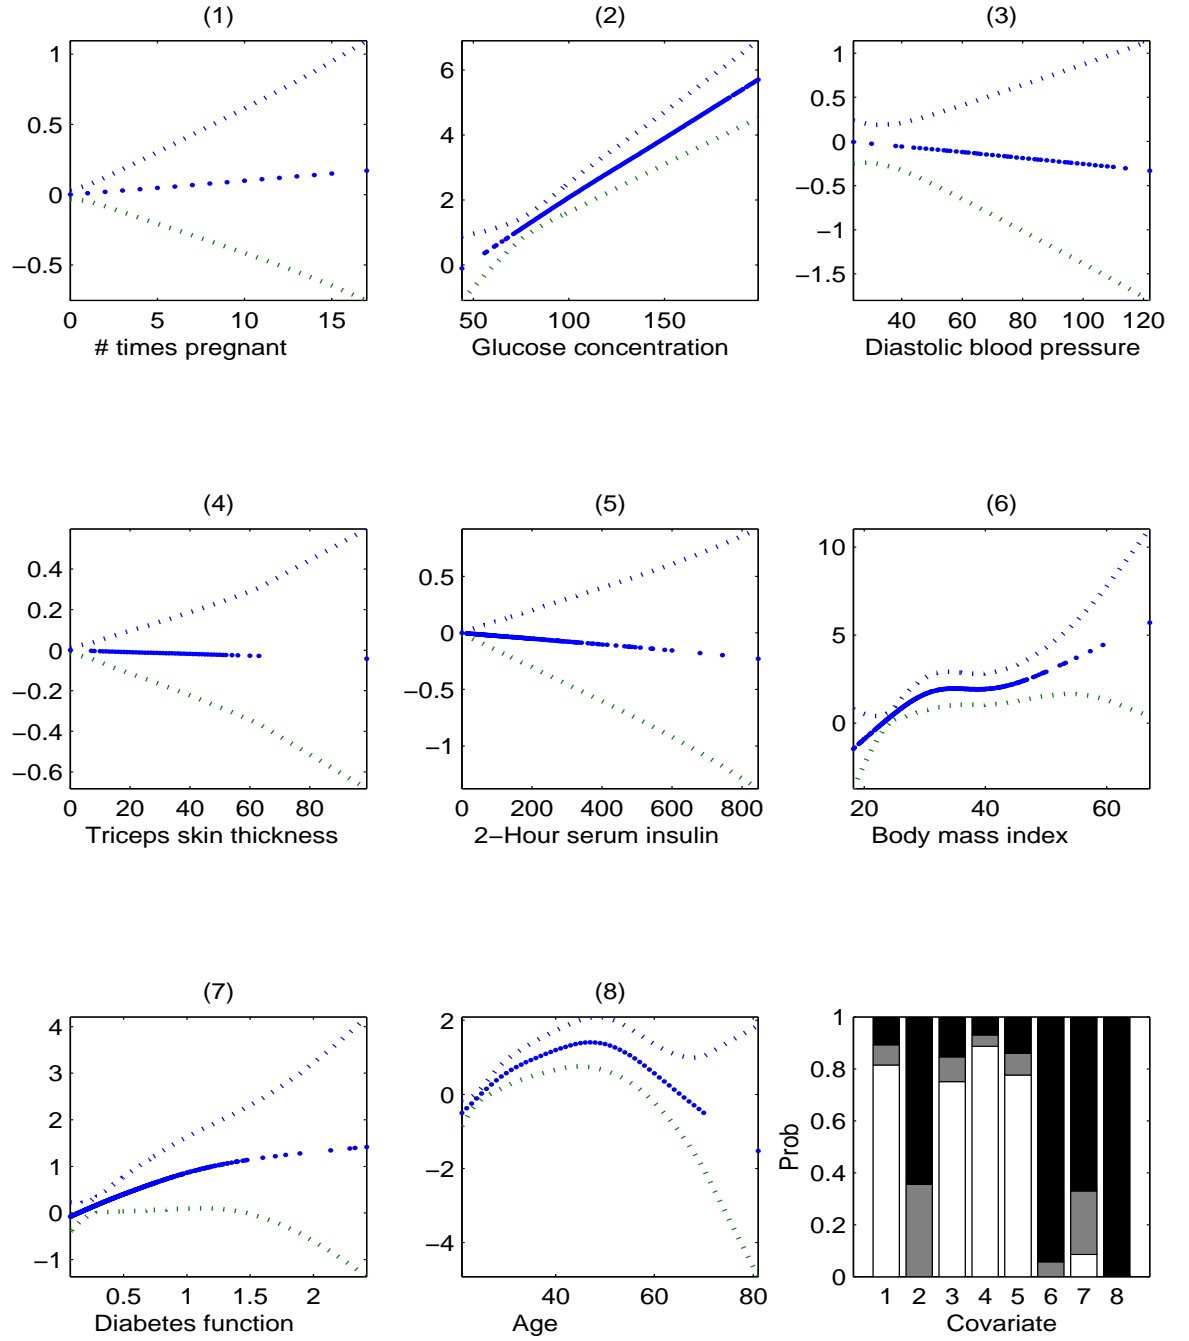


FIGURE 3. Logistic Diabetes data. Plots of the posterior means of the covariate effects at the design points (dotted bold line) and 95% credible intervals (dotted lines). The barplot gives the posterior probability of each covariate function being null (white), linear (grey) and flexible (black).

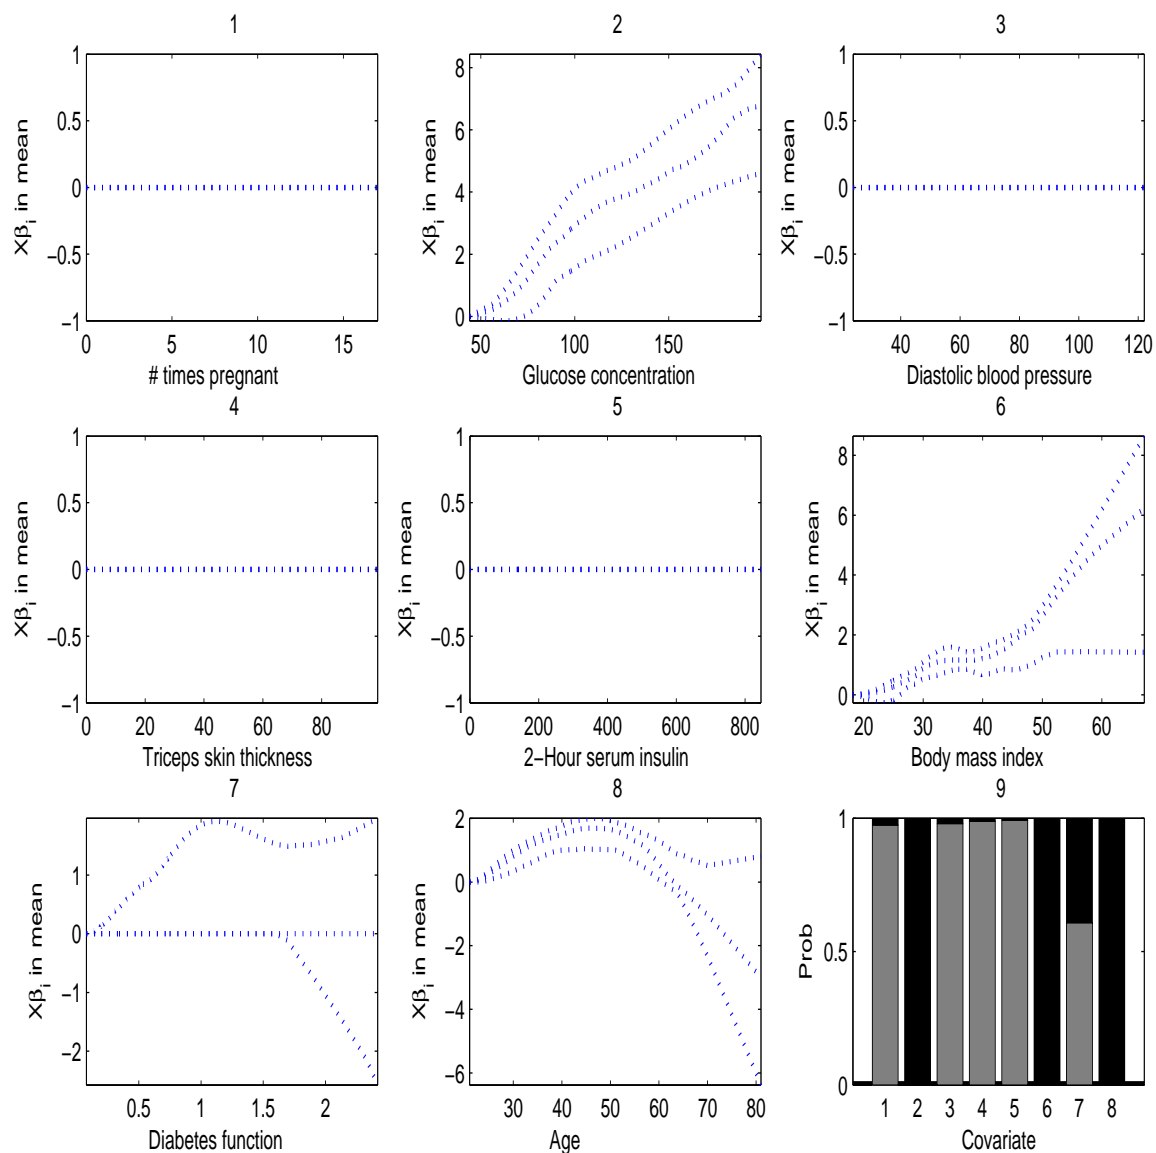


FIGURE 4. Logistic Diabetes data. Plots of the posterior means of the covariate effects at the design points (dotted bold line) and 95% credible intervals (dotted lines) for data based priors approach. The barplot gives the posterior probability of each covariate function being in or out of the model.

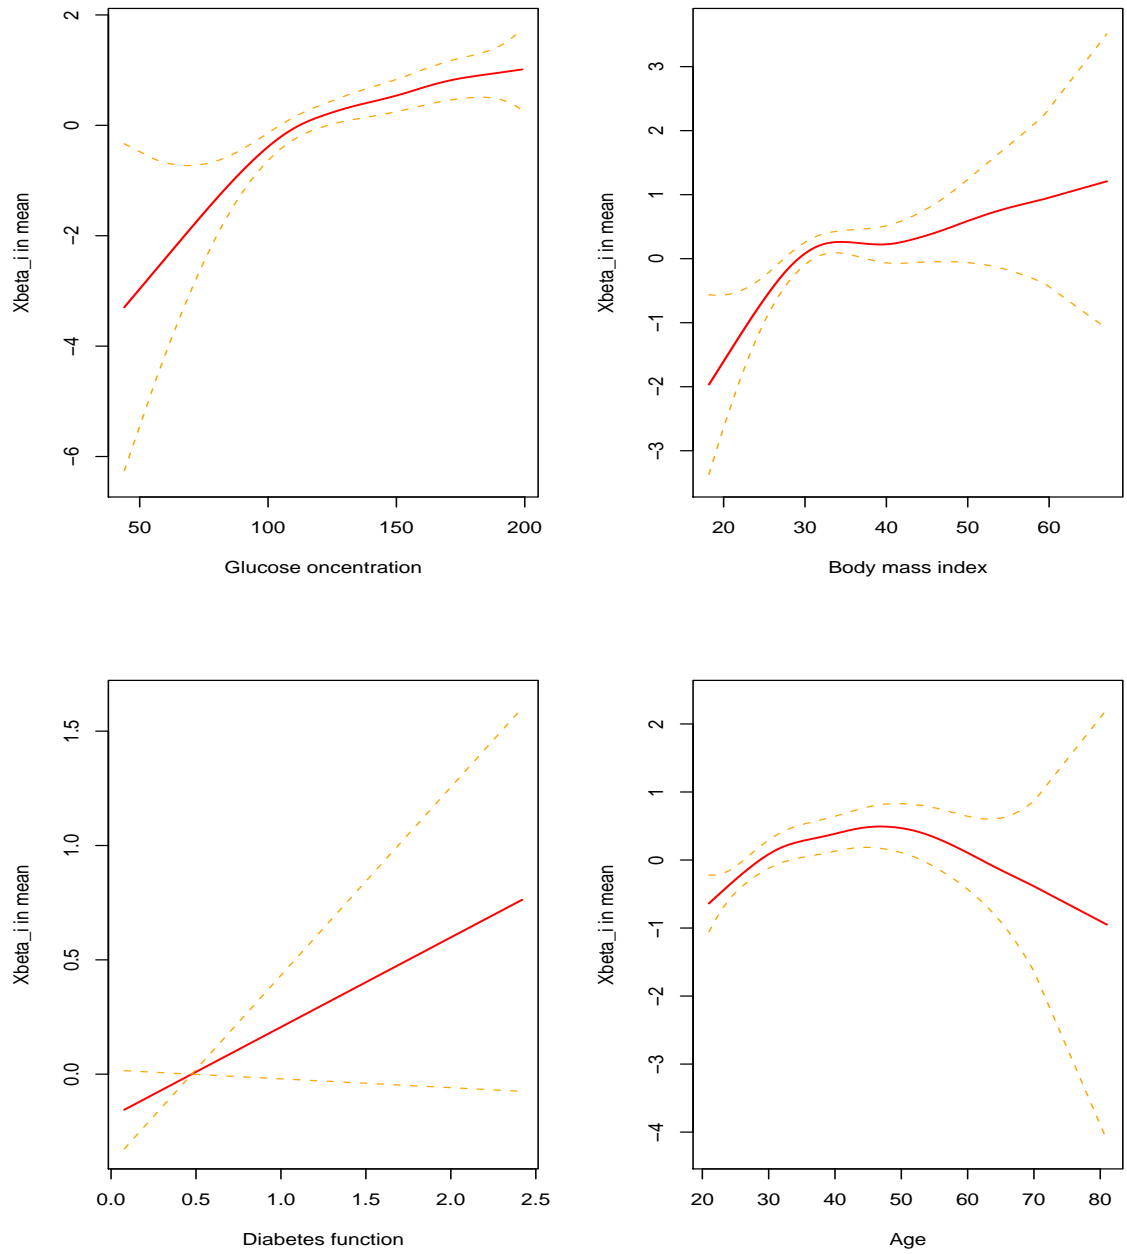


FIGURE 5. Logistic Diabetes data. Plots of the estimated covariate effects (solid line) and 95% credible intervals (dotted lines) using GAMLSS and a backward stepwise model selection procedure with generalized AIC criterion for model selection.

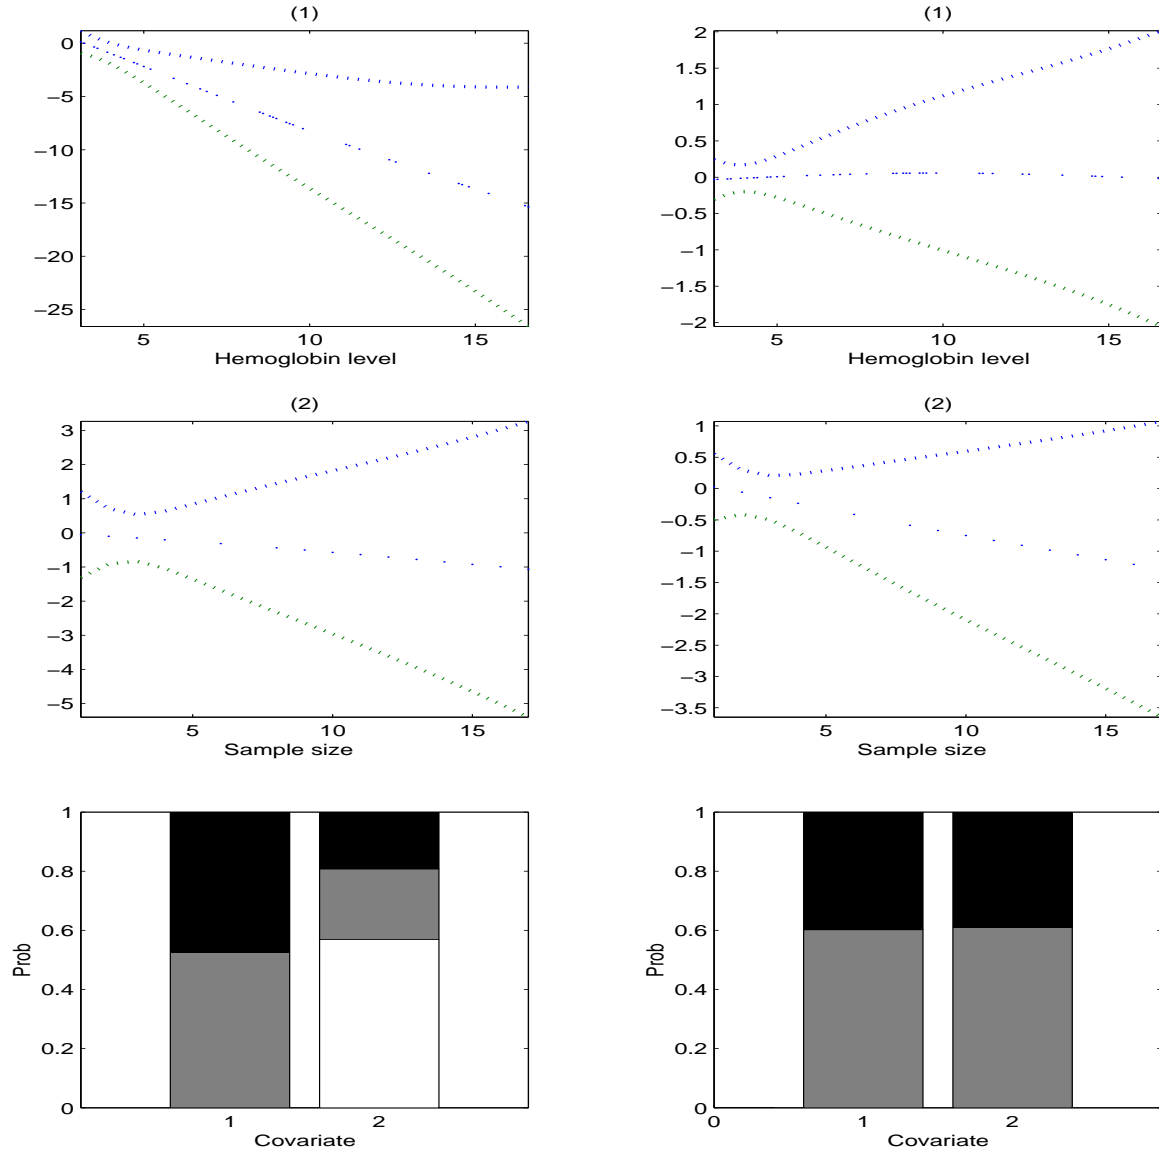


FIGURE 6. Double binomial rats data. Left column: Plot of the posterior means of the effects in the mean model (bold dotted line) together with the 95% credible intervals (dotted lines). The barplot plots the posterior probability of the effects being null (white), linear (grey) and flexible (black). Effects for the dispersion component are plotted in the right column.

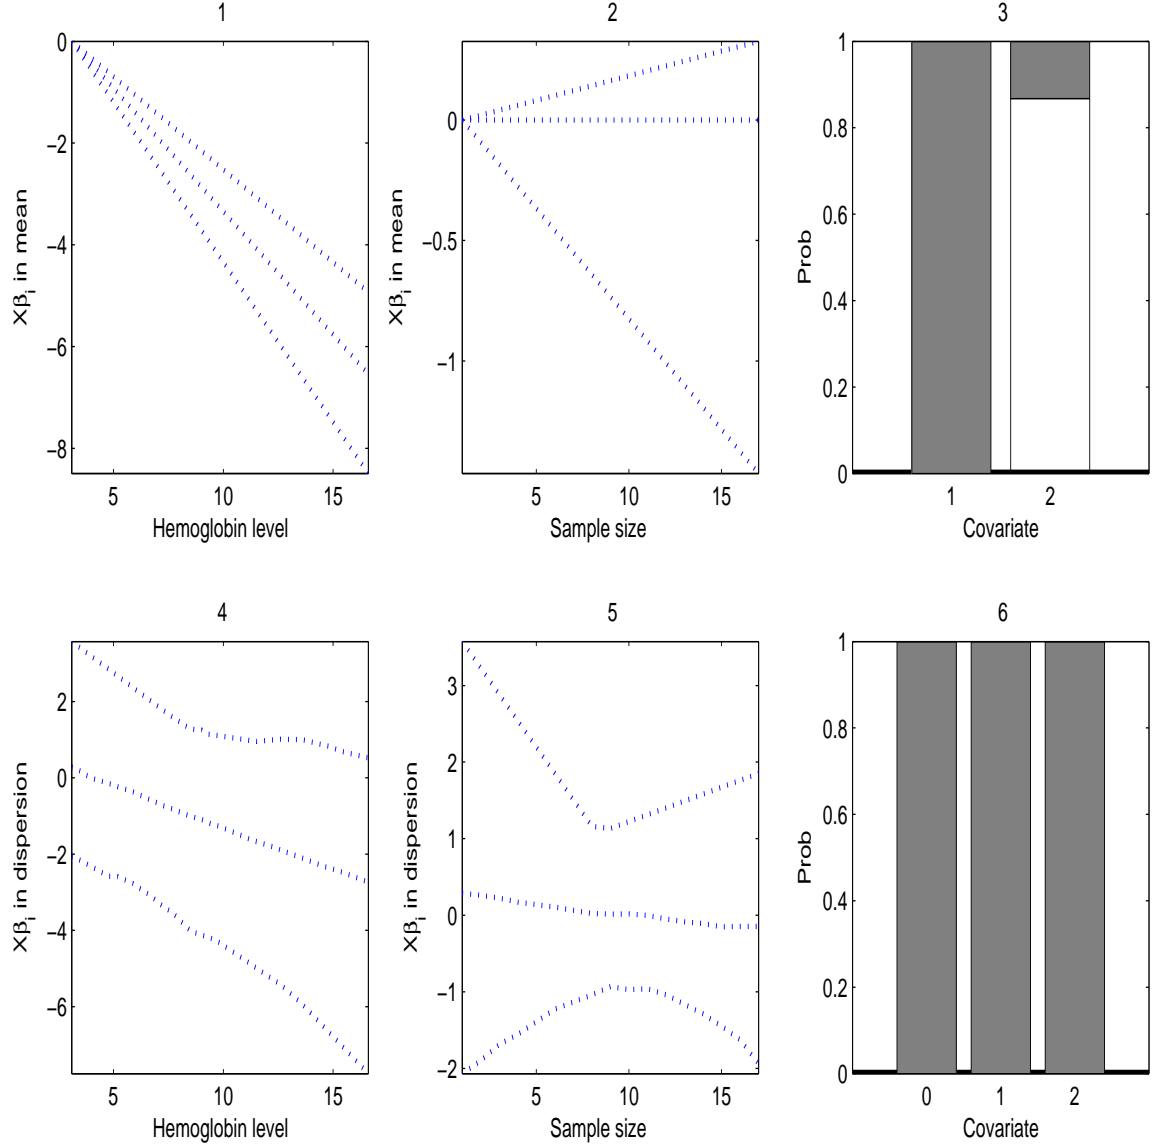


FIGURE 7. Plots of effects for hemoglobin in the mean and dispersion models (left top and bottom respectively) and for sample size (middle top and bottom respectively) together with pointwise 95% credible intervals (dotted lines) for rats data. On the right, the barplots show the probabilities for null (white), linear (grey) and flexible (black) effects.




FIGURE 8. Plots of effects for hemoglobin in the mean model and dispersion models (top and bottom respectively) together with pointwise 95% credible intervals (dotted lines) for rats data where the fit is obtained from the GAMLSS package.